

Textual Entailment

Part 5: Multilingual, Component-based System Building

Sebastian Pado

Institut für Computerlinguistik

Universität Heidelberg, Germany

Rui Wang

Language Technology

DFKI, Saarbrücken, Germany

Tutorial at AAI 2013, Bellevue, WA

Thanks to Ido Dagan for permission to use slide material

Structure of the Tutorial

- Part 1 [SP]: Introduction and Basics
- Part 2 [RW]: Classes of Strategies and Learning
 - * BREAK*
- Part 3 [SP]: Knowledge and Knowledge Acquisition
- Part 4 [SP]: Applications
- **Part 5 [RW]: Multilingual, Component-based System Building**

State of the Art

- What is the state of the TE community in 2013?
 - Almost ten years of research
 - Where do we go from here?
- **Evaluation**: gain insights on what works
- **Sustainable development**: build systems that reflect these insights
- **Application**: make a difference for NLP with TE

3

State of the Art (cont.)

- In MT, there is a “universal platform”
 - MOSES (Koehn et al., 2007)
- There are two open source systems for TE:
 - EDITS, an alignment-based system
 - BIUTEE, a translation-based system
- So people can download these systems, experiment with them, and use them in applications?
 - In principle yes...
 - ...but there are a couple of problems

4

Problems

- Systems are prototypes of specific algorithms
 - Hard-wired preprocessing tools
 - Hard-wired assumptions about language
 - No modularization of algorithmic parts
 - No interchange format for inference rules

In sum:

Evaluation, development, application are difficult

Are we back at square one?

5

Summary

- Theoretically
 - Reusability of Algorithms and Resources
 - Framework Generality
- Practically
 - Systematic Evaluation
 - Multilinguality, and Integration in Applications

6

The EXCITEMENT Project



EXCITEMENT EXploring Customer Interactions through Textual EntailMENT

- EXCITEMENT Open Platform (EOP)
 - Multilingual
 - Component-based
 - Open source
- <http://www.excitement-project.eu>

7

The EXCITEMENT Project


- EU FP 7 Project
 - HEI, DFKI, Bar-Ilan, FBK + industrial partners
- Goal: Provide the necessary infrastructure for sustainable research in Textual Entailment
 - **Specification:** Modular architecture for TE systems
 - Reusability of algorithms, resources through interfaces
 - Towards “plug and play” construction of systems
 - **Platform:** Implementation of modular specification
 - Working for English, German, Italian

Complete

Complete

8

The EOP Architecture



9

Specification

- Linguistic Analysis Pipeline
 - Apache UIMA: linguistic analysis = enrichment of document with strongly typed annotation
 - DKPro type system: language-independent representation of (almost) all linguistic layers
- Entailment Core (Java-based)
 - Interfaces for relevant modules
- Also: “soft” constraints (“best practice” policies)
 - Initialization behavior, error handling, ...

10


Entailment Core

- Top-level interface: Entailment Decision Algorithm
 - Text-Hypothesis pair (UIMA) in, Decision out
 - Existing systems can be wrapped trivially as EDAs
- Three major component types
 - Annotation components
 - Feature components
 - Knowledge components
 - (Don't cover everything, but 95%)

11

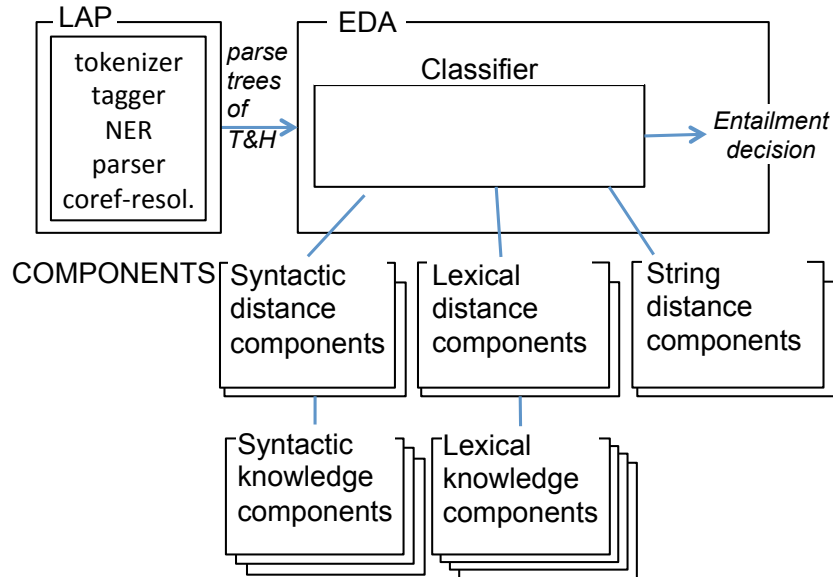
Components

- Annotation components
 - Add linguistic analysis to the P/H pair, e.g. alignment
- Feature components
 - Compute match/mismatch features, distance/similarity features, scoring features, ...
- Knowledge components
 - Provide access to inference rule bases



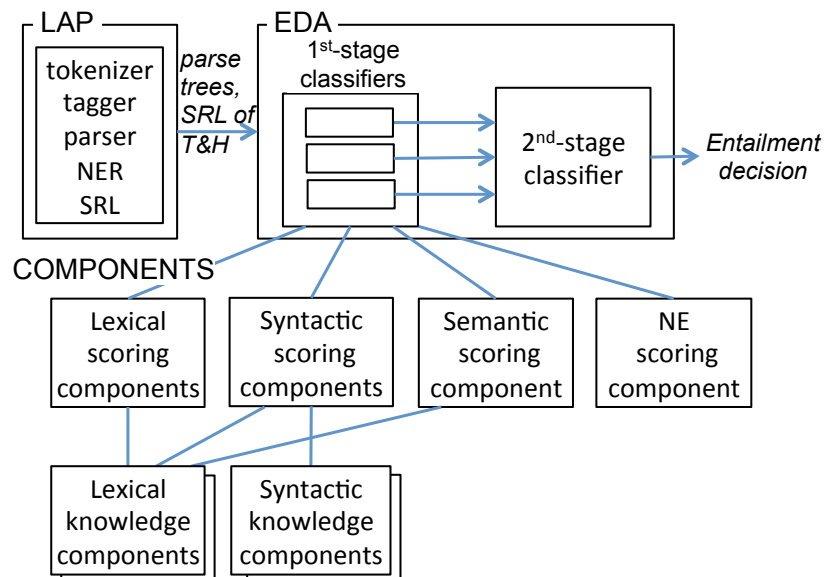
12

EDITS




13

TIE

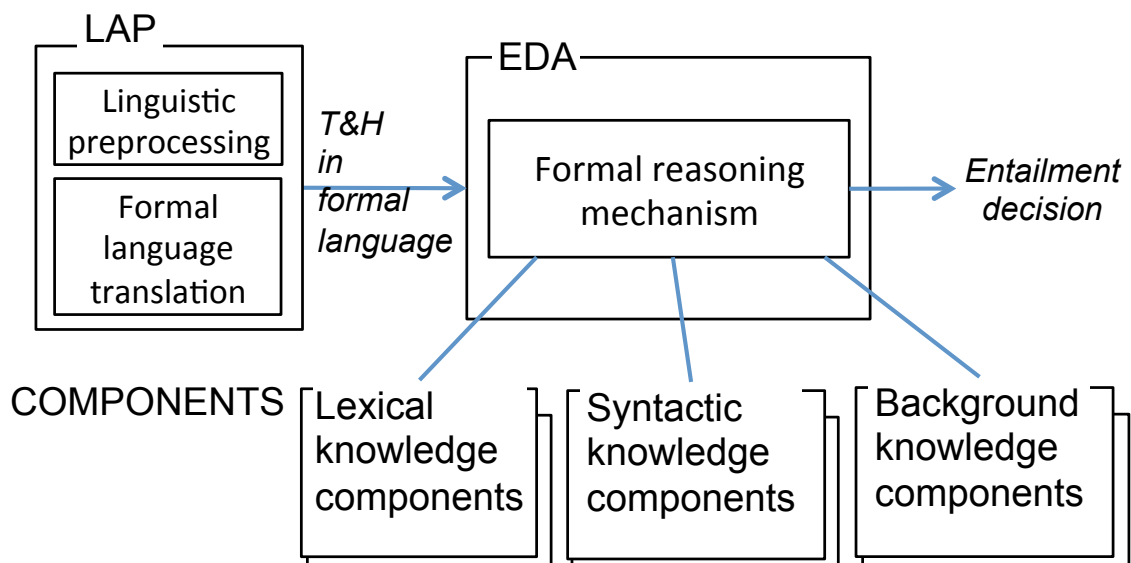


14

BIUTEE



A Formal Reasoning System



Status

- Datasets (Based on RTE-3 data)
 - English, German, Italian, 1600 T-H pairs for each
- LAPs
 - For three languages
- EDAs
 - Three EDAs, EDITS, TIE, and BIUTEE
- Various components
- ...and Many knowledge resources

17

Benefits and further plans

- Reusability
 - Import of BIUTEE's large lexical resources into EDITS for more informed syntactic distance measures
 - Use TIE's semantic role labeller to extend BIUTEE's knowledge resources
 - **“Toolbox” for future experiments**
- Comparable settings for experiments across EDAs
 - constant resources, constant preprocessing, ...
- **Platform will be open-sourced**
 - Community of users

18

System Demo

Subscribe to:

<http://hltfbk.github.io/Excitement-Open-Platform/mail-lists.html>

Public
release on
August 1st!

19

Wrap-Up

20

Structure of the Tutorial

- Part 1 [SP]: Introduction and Basics
- Part 2 [RW]: Classes of Strategies and Learning
- Part 3 [SP]: Knowledge and Knowledge Acquisition
- Part 4 [SP]: Applications
- Part 5 [RW]: Multilingual, Component-based System

De De Explore new application scenarios
re kn • General semantic relation between texts
•
•

21

Not Covered in this Tutorial

- Formal reasoning methods
 - Tatu et al. (2006); Bos and Markert (2005); MacCartney and Manning (2007); Clark and Harrison (2009a,b)
- Corpus construction
 - Cooper et al. (1996); Burger and Ferro (2005); Wang and Sporleder (2010); Wang and Callison-Burch (2010)
- Related tasks: Paraphrase acquisition, Semantic textual similarity, etc.
- Crosslinguality: Mehdad et al. (2010)

22

Further Reference

- Tutorials
 - Dagan et al. ,ACL 2007
 - Sammons et al., NAACL 2010
 - Wang, HIT-MSRA Summer School 2012
 - <http://mitlab.hit.edu.cn/2012summerschool/>
 - Zanzotto, Web Intelligence 2012
 - http://art.uniroma2.it/zanzotto/teaching/tutorials/rte_at_web_intelligence/
- ACL RTE resource pool
 - [http://aclweb.org/aclwiki/index.php?title=Textual Entailment Resource Pool](http://aclweb.org/aclwiki/index.php?title=Textual+Entailment+Resource+Pool)

23

Further Reference

- Book
 - Dagan, I., Roth, D., and Zanzotto, F. M. (2012). Recognizing Textual Entailment: Models and Applications. Number 17 in Synthesis Lectures on Human Language Technologies. Morgan & Claypool.
- Book chapters & Journal Articles
 - Dagan, I., Dolan, B., Magnini, B., and Roth, D. (2009). Recognizing textual entailment: Rational, evaluation and approaches. Natural Language Engineering, 15(4).

24

Further Reference

- Book chapters & Journal Articles
 - Androutsopoulos, I. and Malakasiotis, P. (2010). A Survey of Paraphrasing and Textual Entailment Methods. *Artificial Intelligence Research*, 38:135–187.
 - M. Sammons, V.G. Vydiswaran, and D. Roth (2012). Recognizing Textual Entailment. In: *Multilingual Natural Language Applications: From Theory to Practice*.
 - S. Pado & I. Dagan. (to appear). Textual Entailment. *Oxford Handbook of Natural Language Processing*.

Thank YOU!

Subscribe to:

<http://hltfbk.github.io/Excitement-Open-Platform/mail-lists.html>

Reference List

- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., and Bojar, O., Constantin, A., and Herbst, E. 2007. Moses: Open source toolkit for statistical machine translation. In Proceedings of ACL.
- Tatu, M., and Moldovan, D. 2007. Cogex at RTE3. In Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing.
- Bos, J., and Markert, K. 2005. Recognising textual entailment with logical inference. In Proceedings of HLT-EMNLP.
- MacCartney, B., and Manning, C. D. 2007. Natural logic for textual inference. In Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing.
- Clark, P., and Harrison, P. 2009. Large-scale extraction and use of knowledge from text. In Proceedings of the fifth international conference on Knowledge capture.

Reference List

- Clark, P., and Harrison, P. 2009. An inference-based approach to recognizing entailment. Proc. of TAC.
- Robin Cooper, Dick Crouch, Jan Van Eijck, Chris Fox, Johan Van Genabith, Jan Jaspars, Hans Kamp, David Milward, Manfred Pinkal, Massimo Poesio, and Steve Pulman. 1996. Using the framework. FraCaS Deliverable.
- Burger, J., and Ferro, L. 2005. Generating an entailment corpus from news headlines. In Proceedings of the ACL Workshop on Empirical Modeling of Semantic Equivalence and Entailment.
- Wang, R., and Sporleder, C. 2010. Constructing a textual semantic relation corpus using a discourse treebank. In Proceedings of LREC.
- Wang, R., and Callison-Burch, C. 2010. Cheap facts and counter-facts. In Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk.
- Mehdad, Y., Negri, M., and Federico, M. 2010. Towards cross-lingual textual entailment. In HLT-NAACL.